DOCUMENT RESUME

ED 028 414 AL 001 761

By-Travis, Charles S.

How We Speak and Understand English.

Pub Date [Mar 69]

Note-23p.; Paper presented at the Southeastern Conference on Linguistics, Florida State University, Tallahassee, March 28-30, 1969.

EDRS Price MF-$0.25 HC-$1.25

Descriptors-*English, *Linguistic Competence, Linguistic Performance, Linguistic Theory, *Listening Comprehension, Models, *Speaking, Syntax, *Thought Processes

Identifiers-*Decoding, Encoding

In this paper the author examines the "encoding-decoding" model of speaking and understanding English. He reviews in detail an objection to the model: that it was specifically designed with a view to incorporating linguistic theories, such as syntax, into it. As a result, what it more or or less accurately represents is the relation between the linguist and some language other than his own which he may be studying. However, this relation is necessarily quite different from that holding between the competent speaker-hearer and his own language. What seems to be required of an explanation of how we understand English is that it describe how, on hearing an utterance, we attain an appropriate terminal state. This is a description that the encoding-decoding model clearly fails to provide. What it does describe is a terminal event (the emission of a signal). Even if the model were revised so that it did describe a terminal state, that state could not account for understanding an utterance. The author believes, however, that (with a few refinements in the model) there is no reason to take the objections seriously at all. He next discusses the encoding part of the model which accounts for how we speak English. (DO)

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
OFFICE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL OFFICE OF EDUCATION POSITION OR POLICY.

ED028414

How We Speak and Understand English

A competent speaker of English is capable of understanding at least some, perhaps much of what is said in it. He is also often capable of expressing himself in it. That is, generally speaking, he can at will produce an utterance that is to be understood in English as what he wants to say. If he wants to say that there are spotted toadstools growing under the azaleas, then he can produce an utterance that will be so understood. In most, but not all contexts, "There are spotted toadstools under the azaleas" will do. The puzzling question is, how does he do all this? In this paper I will show what sort of an answer that question requires.

A partial answer to the question of how we speak and understand English is that we know many facts about it. Ignoring for the moment questions about individual variation, we can say exactly what relevant facts we know about English by providing an adequate English syntax, together with whatever else may be necessary for a complete theory of linguistic description.

Such an answer is not likely to satisfy us. It is an answer of a sort, but there is much that comes to mind that is left unanswered. In particular it seems natural to ask how we make use of our knowledge in understanding or producing specimens of English.

One very natural answer to this question is sketched, in the case of "The Speaker of a Language" by Jerrold Katz.[1] Following

---

1. "Mentalism in Linguistics", Language, 1964.

Gilbert Harman. I shall call this the encoding-decoding model.

Katz says:

> The speaker, for reasons that are biographically but not linguistically relevant, chooses some message he wants to convey to the hearer. He selects some thought he wishes to express to him, some command he wants to give him, or some question he needs to ask him. This message is, we may assume, in whatever form the semantic component of his linguistic commands, questions, or the like. The speaker then uses the sentence production procedure to obtain an abstract syntactic structure having the proper conceptualization of this thought. . . .
>
> Although this model is phrased as if the processes described were conscious, no such assumption is involved. It is not an essential feature of mentalism that the processes postulated by the mentalist as going on inside the speaker's head should be open to the speaker's conscious awareness.

There is a fairly standard sort of objection to be made to the encoding-decoding model. It is an objection that has been raised on several occasions by Harman,[2] and at one time by me.[3] I will now review this objection in some detail. I will then show that the objection is fundamentally mistaken and that there is much right with the encoding-decoding model.

II

The objection to the encoding-decoding model may be summed up as follows. The model was specifically designed with a view to incorporating linguistic theories, such as syntax, into it. As a result, what it more or less accurately represents is the relation between the linguist and some language other than his

---

2. "Psychological Aspects of the Theory of Syntax", Journal of Philosophy, February, 1967, and Language, Thought and Communication, unpublished.

3. Cf. my dissertation, Innate Ideas, 1967.

own which he may be studying. However, this relation is necessarily quite different from that holding between the competent speaker-hearer and his own language.

If a linguist has constructed a theory which provides adequate descriptions of some language, then it seems plausible that, with more or less reliance on his intuition, he can do the following things. He can, at a minimum, take any sentence in the language and provide it with a syntactic description and semantic interpretation either in some other language or in some universal semantic code. Conversely, he could take any "message" expressed in some universal semantic code and find a sentence in the language which conveyed it. With a much stronger theory, and perhaps more reliance on intuition, he could take any regular well-formed utterance in the language, describe it syntactically and say what it meant in the vocabulary of some universal theory. Conversely, he could take any thought and find a contextually appropriate way of expressing it in the language.

One goal that linguistics could set for itself is to find mechanical procedures for using theories of linguistic descriptions to get from thoughts to their expression in a language, or vice-versa, so that the linguists' reliance on his intuitions was reduced to an absolute minimum.[4] If this enterprise were successfully carried out, linguists would have succeeded in creating an artificial

---

4. This, of course, is not a goal of generative grammar, as presently conceived, nor is it a goal linguists are likely to adopt. It is brought in here only to show the relation of the encoding-decoding model to generative theories.

linguistic competence.[5] That is, anyone who could use the procedures could, with no knowledge of the language in question, converse in it. (It is irrelevant here that the complexity of the procedures may make conversation painfully slow.) This artificial linguistic competence would be an adequate realization of the encoding-decoding model. The interesting question, however, is whether this artificial competence also describes the sort of linguistic competence a normal English speaker has. There is good reason for saying that it does not.

To bring out the disparity between this artifical competence and the real thing, let us first look at how an English speaker understands what is said in English. What is required by the present model seems to be something like the following. First, speech sounds are operated on by some preliminary analyzing devises. These devices identify the sounds as candidates for further analysis by the hearer's linguistic mechanisms. The sounds, thus identified, are then processed by further devices which, at a minimum, provide them with the correct syntactic analysis and semantic interpretation, together with whatever other sort of linguistic interpretation of the utterance is required. This information is presumably in the form of some universal semantic code. The devices in question make use of the hearer's knowledge of English in the way indicated by the procedures of the linguist's artificial linguistic competence. Let us call them the decoder. When the decoder outputs the information just described, caeteris paribus, the hearer understands what was said.

---

5. Of course, this is quite different from describing a conversational robot. No claims will be made in this paper about that sort of creature.

It is natural to ask at this point what it is like for the decoder to output a message. It is difficult to find any interpretation of this abstract description on which the process might be helpful to the hearer and it is not patently false that it occurs. On one natural interpretation, the message the hearer receives is simply what was said. For example, if Sam says to Pete, "It's going to rain on Saturday", then Pete receives the message that it's going to rain on Saturday, expressed as "It's going to rain on Saturday". There are two difficulties with this interpretation. First, on it the message is not produces by the linguistic analyzing mechanism but, in the above case, for example, by Sam. Second, we still lack an explanation of how it is that Pete understands what Sam said. For what we wanted to know in the first place is, how did Pete understand Sam when Sam said "It's going to rain on Saturday." We started out wanting to know what special properties Pete has, in virtue of which he can do what not every human being, much less every conceivable organism can, that is, understand utterances in English. What our answer amounts to on this interpretation is that he understands English. Clearly, then, progress does not lie in this direction.

On any other interpretation of outputting a message, there are two difficulties. First, it seems patently false that the hearer ever recieves a message when he understands what someone said. Second, the process we have described still fails to explain how anyone could understand English.

Normally, when an English speaker perceives an English utterance and understands it, he doesn't perceive anything besides the utterance.

He is not aware, either as something he thinks of, or as something he hears or sees, of any message expressed in the vocabulary of any theory of linguistic descriptions. Moreover, most English speakers have not mastered such a vocabulary. Being aware of such a message would not help them understand the utterance in any case.

Even if English speakers had mastered the appropriate messages, the explanation in terms of the linguist's artificial linguistic competence would still fail. For given that not all people, much less all conceivable organisms, can understand messages expressed in the vocabulary of a given semantic theory, we will want to know what special properties English speakers possess in virtue of which they can understand such messages. Since there are at least as many messages as there are ways of understanding utterances in English, and at least as many of them are novel for any given English speaker as there are utterances in English which are novel for him, this would seem to be a problem of roughly the same complexity as the original.

The present situation is a curious one. We begin by asking how English speakers understand what is said in English. The answer suggested by the encoding-decoding model is that they possess certain devices which "tell" them what English utterances mean. But then we ask how they understand what these devices "tell" them. If the encoding-decoding model is compelling at all, then it would seem to be compelling here as well. The answer would thus be that English speakers possess further devices which enable them to understand the output of their linguistic analyzing mechanisms. But, first, that is

no answer at all. Second, it involves us in a vicious infinite regress.

Although the objection has been phrased so far in terms of the English speaker as a hearer of utterances, it applies just as well to the problem of how he produces what he wants to say. According to Katz, the speaker begins by choosing a message. But how does he know that any message he chooses expresses what he wants to say? How does he even distinguish between one message and another? Given that there are mechanical procedures for producing sentences conveying a given message, what procedures could there be for producing messages expressing the right thought?

III

The difficulties that have appeared in the encoding-decoding model suggest that it imposes the wrong form on an answer to the question how we speak and understand English. A look at the concept of understanding also indicates that this is the case. First, understanding is a state. Hence we cannot say what it is to understand by describing a process whose end product is an event. Second, on any account the decoding process does not terminate in a state of the right sort.

Saying that understanding is a state means that when someone understands an utterance an indefinite number of his potential performances are modified because of it. That is, he is in a position to do an indefinite number of things he could not otherwise do. Some of the performances characteristically associated in

this way with understanding an utterance are, repeating it, paraphrasing it, drawing inferences from it, carrying out instructions contained in it and modifying strategies and attitudes because of it. For example, if Sam says to Pete, "Please shut the door", then if Pete understands, he can say that what Sam said was "Please shut the door". He can also say that Sam asked him to close the door. He can shut the door, or refuse to do so, so that his actions accord with his attitude towards Sam. He could, under certain circumstances, form the opinion that Sam is pushy or demanding, using Sam's request as a reason. He might also adopt the policy of trying to avoid Sam when he was tired.

Of course, there is no fixed performance that we require of a person if we are to say that he understands a given utterance, nor is there any fixed quantity of performances that he must produce. There is room for indefinite variation among individuals in their responses to almost any utterance. The performances we expect of a given individual on a given occasion depend very much on what we know about the person, as well as on our interests in asking whether he understands. If Sam says "There is fresh blueberry pie at Ferguson's Cafe", a sign that Pete understands may be that he immediately goes to his bookie to collect on his bet on the world series, if Pete believes that Ferguson serves fresh blueberry pie only on special occasions such as when his team, which Pete knows is also Pete's team, wins the series. But if a person fails in enough of the performances we would normally expect of a competent English speaker who did understand, then we will say that he does

not understand what was said.

What seems to be required of an explanation of how we understand English, then, is that it describes how, on hearing an utterance, we attain an appropriate terminal state. That is, we want to know how the speaker gets into a position to be capable of the performances he could not otherwise produce, and what it is for him to be in that position. This is a description that the encoding-decoding model clearly fails to provide. First, as it stands, it doesn't provide a description of a terminal state at all. What it describes is a terminal event, that is, the emission of a signal. Second, even if the model were revised so that it did describe a terminal state, that state could not account for understanding an utterance.

The model could be revised in a number of ways so that the decoding process would be described as a series of changes of state in some device. No matter how this was done, a description of the terminal state would contain, at most, the same information contained in a description of the terminal signal on the unrevised model. To decode, in this case, is to go from a description of an acoustic event to a description of something said in English. The latter is precisely what is contained in the output signal on the unrevised model. Any additional information contained in a description of the terminal state is not information one could extract from the utterance simply by knowing what is said by it in English. Showing how such information is attached is not, then, modeling competence in English.

A description of a terminal state which contained no more information than that described above could not completely account for our understanding an utterance. When we understand an utterance

we modify an indefinite number of potential performances. These performances are not necessarily themselves linguistic. They need not have to do with what we say about the utterance or about the language. But a description of the terminal state of the decoding device does not explain how we modify any performances at all, much less ones that are not themselves exercises of our competence in English. In short, showing how we acquire information is not the same as showing how we are in a position to use it as we do.

Understanding being what it is then, it is easy to see how any decoder may fail to account for it. We can always describe someone's failure to understand in terms of his failure to produce enough of the right kind of performances. Caeteris paribus, such a description will defeat any claim that the person does understand. But such a description is logically compatible with the production within the speaker of every sort of information about what was said. So it might happen that while a person's decoder works perfectly he consistently fails to understand anything that was said.

It is also easy to see how objections like Harman's can always be generated for any decoder. We need only point to the above kind of failure to understand and ask why sometimes, at least, it doesn't happen. Clearly the decoder doesn't explain this. But then one could claim that the decoder can't explain understanding either, for if it is any part of such an explanation then there must be another part to show why decoders do help people get into the right states.

Similar remarks pertain to the speaker. Frequently in speaking English the speaker produces utterances which express what he intended to say. We cannot account for this by postulating the occurrence of a

signal, or thought, and then showing how a particular utterance was matched up with it. We will always need to explain how the speaker produced a signal such that the utterance which matches up with it expresses what he wanted to say. Again, having the intention of saying something is being in a state. We need an account of how speakers come to be in such a state, and how being in it results in the production of an appropriate signal.

IV

Having seen the underlying reason why we can always generate objections to the encoding-decoding model like those expressed in section two we can now ask just how fatal they are. I will suggest in this section that, with a few refinements in the model, there is no reason to take the objections seriously at all.

The first refinement that must be made concerns the output signal. This signal cannot be equivalent to a semantic "message" or "reading". Such messages express the meaning of sentences, not thoughts, and not what is normally expressed in producing an English utterance. For example, different utterances of "John went to the bank" are rarely to be understood in the same way, even on a given reading of "bank". There are many Johns and many banks for the utterance to be about, and the utterance will generally be understood differently where the references of these terms differ. The thought that no particular John went to no particular bank is hardly a thought at all. The thought that some contextually definite John went to some contextually definite bank is not the thought normally

went to some contextually definite bank is not the thought normally expressed by "John went to the bank". Elsewhere[6], I have introduced the notion of an illocutionary structure, which is a description containing the extra information necessary to say how a given utterance in a given language is to be understood. Let us suppose that the output of the decoder is characteristically one or more illocutionary structures.

Second, if (as Katz suggests) the notions of a speaker choosing or receiving a message are not to be taken seriously, then we ought to say what their serious equivalents are. Let us suppose that the mental signals which occur in the encoding-decoding model are neural events. On the model, when a speaker understands an utterance the decoder produces a certain signal. Thus, in this case, the model is of a speaker in whom a neural event occurs fitting that description. Similarly, the decoder is to be taken as some sort of neural mechanism. For someone to understand an utterance on a given occasion is for him to enter some appropriate state, brought about by the production of the appropriate signal by his decoder. The explanation provided by the encoding-decoding model is, then, that the decoder, on receipt of an acoustic signal, produces a signal consisting of one or more illocutionary structures in a way made explicit by the model. These structures then cause the hearer to enter an appropriate state.

Before asking whether, in the light of the objections to it, this explanation is acceptable, we ought to ask why anyone should ever think that such a process of decoding is a necessary part of an

---

6. "A Generative Theory of Illocutions", dittoed.

explanation of understanding an utterance. The answer is that such a process is necessary if the hearer is to distinguish what he has heard from everything else that might have been said in English. The information contained in the signal produced by the decoding is precisely what is needed to know what was said. The hearer cannot treat this utterance differently than he would have treated a large number of other utterances unless his treatment of the utterance is determined by all of the information contained in the signal.

When the hearer enters an appropriate state, caeteris paribus, we say he understands what was said. But we only say that if he is in the state because he knows what the utterance is and understands English. He might have attained the same state by misperceiving the utterance and entering a state which is inappropriate for what he took the utterance to be. If someone says, "Close the door" I may think he said "I froze the floor" and close the door. But that is not understanding what was said. What determines that this is not what he was doing is that the state he entered was determined by the information contained in the decoding signal. Thus, if the hearer does take what he heard to be "Close the door" and not "I froze the floor", then there must be sufficient information in the signal to provide a description of one sentence and not the other. If he takes the utterance to be about the front door and not the bedroom door, then the signal must have provided a description which determined that the utterance was to be understood in one way and not the other. Similar remarks hold for the other features which determine how the utterance is understood.

Of course, the encoding-decoding model is of an ideal speaker-

hearer who understands what was said perfectly. We often say that real speaker-hearers understand what was said perfectly well when they have access to less information about it than what would be contained in the appropriate illocutionary structure. Sometimes I may even understand what you said without knowing which words you used. But this in no way affects the above argument. It merely means that in real speaker-hearers the signal provided by the decoder is sometimes less rich than described above. But to whatever extent a person does understand what was said, to that extent his relevant state must be determined by the information provided by the decoder on the encoding-decoding model.

Now the objection to the encoding-decoding model is that it fails to explain how the hearer gets into an appropriate terminal state. We cannot deduce from the descriptions given by the model that the speaker does enter such a state. We can conceive of his undergoing all the processes described in the model and not doing so.

If this is all that is left of the objection, then it is clearly invalid. The requirement it places on explanations is simply too high. We don't normally require that an explanation provide us with conditions under which it is logically necessary that the result occurs. An explanation of how an adding machine adds need not explain why it does not turn to dust instead.

Whether the objection is valid, then, ought to depend on whether the information not supplied by the model is information that reasonably should have been supplied. I will argue that there are two reasons why such information should not be required. First, it is not necessary. Second, it is impossible to provide it within any theory of competence in speaking English.

The information is not necessary because what we are modelling is linguistic competence in human beings. Human beings are the sorts of organisms that are affected in particular ways by neural signals of a given sort. English speakers need not make use of their knowledge of English for the signals to have the effects they do. An explanation of how people use what they know about English to speak and understand it need not explain why the signals act as they do.

A theory about the linguistic competence of English speakers could not provide the missing information because, as we have seen, the states of understanding that people attain are indefinitely varied. If two people understand an utterance in the same way, that is sufficient reason to suppose that their decoders produced signals of the same form. But there is no reason to suppose that a given signal ever has the same effect on two different people. If it does, that is not a fact about linguistic competence, but a fact about certain people's knowledge, beliefs, strategies and intentions concerning what is being talked about. This is a point that is made clear by an earlier example. If an English speaker understands "There's fresh blueberry pie at Ferguson's cafe", will this affect his potential performances with respect to a certain bookie? There is no way of telling by studying linguistic competence. Whether the relevant signal from the decoder has that effect on him depends on how he is prepared to receive it.

There is a grain of truth in Harman's position. The performances a person is capable of producing in virtue of understanding an utterance are not accounted for by decoding. If we want to understand why Pete closed the door in virtue of understanding Sam say "Shut the door", we will need to know much more than how a decoder attached information to

a signal. So we might say that the fact that understanding takes a particular form for a particular hearer is not accounted for by decoding. But then, not surprisingly, it is not accounted for by linguistic competence either.

Despite what the decoding model can't do, it is still the strongest available acceptable answer to the question, "how do we speak and understand English?". For some purposes, at least, to ask that question is to ask for the relevant differences between English speakers and non-English speakers in virtue of which the former but not the latter can understand what is said in English. As long as we restrict our attention to human beings possessing a decoder of a certain sort is the only relevant difference there is.

V

Harman has suggested[7] that the encoding-decoding model will seem natural only as long as we concentrate on the role of language in communication. He proposes that if we pay proper attention to the fact that typically one can also think in a language one speaks, the model will not seem compelling. There appear to be at least two reasons why Harman thinks this. First, it is unlikely that one needs to decode one's own thoughts. Second, if we can think in a language then the kind of decoding needed to understand an utterance seems much simpler. We understand an utterance when we take it as expressing the thought that we would think in thinking those words. We understand an utterance of "Snow is white" when we take it to express the thought that we would typically think in thinking the

---

7. "Language, Thought and Communication".

words "Snow is white".

Harman overlooks the complex problems involved both in understanding how we think in a language and in understanding how we take an utterance to express what we would typically think in thinking it. These problems are much the same as we have already encountered in trying to understand how we understand utterances. When we see what they are, we can see that Harman's model fails to avoid the need for the complex decoding we saw to be necessary in understanding what is said.

Assuming that Harman's account of linguistic competence is correct, how do we manage to take an utterance as expressing the thought we might think in thinking it? Why do we take an utterance of "The tulips are blooming" to express the thought that the tulips are blooming, which we might think in thinking "The tulips are blooming"? Why don't we take it to express what we would think in thinking "The girl's lips need grooming"? Part of the answer is that we recognize that these are two different English sentences. But that implies that we can discriminate the sentence uttered from any other English sentence. And that is only possible if the way we take the utterance is determined, inter alia, by all the information contained in its syntactic description. Similar remarks apply to all the other information needed to determine a thought. For example, why do we take an utterance of "The tulips are blooming", to be about some particular bunch of tulips? In short, even on Harman's model, if we understand an utterance then how we take it must be determined by all the information supplied by the decoder on the encoding-decoding model.

Similar problems confront any account of thinking in English. Why is it that in thinking a particular English expression on a certain occasion, we are thinking a particular thought? Why is it that we think the words "The tulips are blooming", when we have the thought that the tulips are blooming? Why don't we have the thought and think different words, as we might if we were thinking in a code? Or why don't we think the words and not have the thought? We might have some other thought instead, or none at all. We might even be puzzled that we thought those words and wonder what they could mean. If it is a fact, then why is it a fact that typically we don't do such things?

An answer to this problem might be that in thinking certain words we have a certain thought in virtue of the way we are prepared to treat our act of thinking. But the situation here is no different from what we found to obtain for understanding an utterance. One can only be prepared to treat an act of thinking in English as an act of thinking one appropriate thought if one's treatment of it is determined by the information contained in the illocutionary structure which correctly describes the corresponding utterance. So, although we don't typically decode our own thoughts, thinking in English may be much closer to talking to oneself than Harman supposes.

VI

Although the decoding model is correct it is easy to be over-optomistic about the extent to which it can be filled in through a study of linguistic competence. Just as individual variation restricts the scope of a theory of understanding utterances, it also restricts

the scope of a theory of decoding them. Such a theory would not be a detailed picture of a decoder, but rather a set of procedures to be used in carrying out specific steps in the decoding. A study of linguistic competence will not reveal how the results of such procedures are used. So such a study doesn't reveal everything about why a decoder produces the results it does. All this can be made clear through an example.

Let us suppose that Sam says to Pete "Go to the store". One thing Pete's decoder must do is to specify the sentence that was uttered. It might do part of this by using procedures available to all English speakers. Such procedures would show that the utterance could be a good instance of the sentence "Go to the store"; or a not quite so good instance of "Go to this door", or a grammatically deviant version of "I'll go to the store", or "I went to the store". (via "I go to the store")

In fact, Pete will understand the utterance as an utterance of one of these. So the decoder adopts just one of these descriptions. Which one it adopts, however, depends on certain assumptions incorporated into it at the time. It will depend, <u>inter alia</u>, on Pete's beliefs about Sam;s normal accent, the relation between Sam's normal accent and the accent he is using on this occasion, Sam's mastery of grammar, possession of his faculties, and so on. What Pete believes on these matters depends partly on his knowledge of Sam, partly on other things, such as whether or not he is paranoid, but not at all on his linguistic competence.

Similar remarks hold for every other feature the decoder must attach. For example, following what is now a standard way of distinguishing speech acts, let us assume that "Go to the store" may be either

an order or a request. One of the conditions for ordering is that the orderer must be in a position of authority over the person he is ordering. One of the sub-routines the decoder might use in identifying this speech act, then, is to ask whether the speaker is in a proper position of authority. But suppose he is not. That will not determine that the utterance is a request. It may be an attempted order that in some respect fails. What determines how Pete will take it? Again, not linguistic competence.

Conflicts between different sub-routines may arise, for example, where syntax indicates that the utterance "Go to the store" is an imperative, but external considerations, e.g., about the discourse, or about Sam, indicate that it was a statement. Part of what the decoder must do in this case is to weigh syntactic against external considerations, and choose one as decisive for the final description. Again, what it does is not determined by linguistic competence. Given certain beliefs about Sam, Pete may take the utterance as an intentionally ungrammatical utterance of "I went to the store". Given others, Pete will take the utterance as an order injected in a pecular way into a discourse.

If understanding involves so much besides linguistic competence, it might seem surprising that we all manage to do it. Two considerations may reduce the surprise considerably. First, sometimes we don't understand what was said correctly. It would be odd to always attribute such failures to defects in linguistic competence. Second, most of us share as hearers a substantial body of knowledge and beliefs which, while not about English, may be used in determining what someone said. Further, as speakers, we share certain regularities in behavior.

There are limits to the deviancies one normally encounters.

VII

I have argued that the encoding-decoding model provides an adequate and a necessary explanation of how we understand English. We have not yet seen whether the model is also appropriate for explaining how we speak it. There is at least some reason for supposing that it is not. On closer examination, however, we can see that speaking does involve encoding.

What would a refined version of Katz's account of encoding look like? Let us take it to be the following. When, for whatever reasons, a speaker is in a position such that he is going to express a thought (or ask a question, give a command, etc.), there is produced in him a signal containing all the information necessary to determine just which thought (or question or command) it is. This signal activates a device which we will call the encoder. The encoder makes use of all of the speaker's relevant knowledge of English to select a syntactic structure describing a sentence which expresses what the speaker wants to say. It then produces a signal which determines that the speaker utters the appropriate sounds. The process of encoding, then, converts a signal containing information which determines a thought into a signal which contains information determining a sentence, and hence an utterance of it.

At first glance, such a process of encoding seems unnecessary. Neither the speaker nor his neural apparatus need determine that any sentence provides the proper expression for any thought. All the speaker need do is utter an appropriate sentence in virtue of being in

a state in which the sentence expresses what he intends to say. Why, then, can't the utterance be determined directly by the signal containing the information determining the thought? Why think that any process of applying what the speaker knows about English is involved in speaking it?

The answer to these questions becomes apparent when we look at the relation between thoughts and sentences. Just as there are often many thoughts that could be expressed by one sentence, there are often many sentences that could, in a given context, express a particular thought. It may often be the case, then, that the sentence produced by the speaker is not fully determined by the thought expressed. It may, to some extent, be determined arbitrarily, or by some set of stylistic maxims that the speaker, on that occasion, adopts. Whether one says "The eye doctor finished with his last patient and left for lunch", or "The oculist, having finished his last consultation, left for his mid-day repast", depends more on one's style than on what one is thinking. Generally, then, the signals that determine thoughts determine equivalence classes of sentences, not sentences. What falls within the class is determined both by the thought and by what means of expressing it syntax allows. The function of the encoder is to use what the speaker knows about syntax, subject to whatever stylistic constraints he adopts, to select one sentence falling within this equivalence class.

The fact that a thought does not always determine a particular syntactic means of expressing it raises an interesting problem. How much of the information in an illocutionary structure, which specifies what was said on a given occasion, may be left out of the specification of a thought?

This problem is difficult because there is much in an illocutionary structure that helps determine syntactic form besides explicit syntactic information. For example, a specification of the references made in a particular utterance, and the relations between them would seem to limit choices of syntactic deep structures. Yet specification of references is often part of specifying a thought. If we could solve this problem, we could see at least two things more clearly. First, we could see what a proposition might be. Second, we could have a better idea of the extent to which syntactic structures are determined by the structure of human thought. But these problems we will not attempt to solve here.

Charles S. Travis

University of North Carolina
at Chapel Hill